

Guaranteeing Termination of Chandrasekaran and Ipsen's Algorithm for Computing Rank-Revealing QR Factorizations*

Gregorio Quintana-Ortí[†] Enrique S. Quintana-Ortí[‡]

Argonne Preprint MCS-P564-0196

ABSTRACT. Many problems from science and engineering require the computation of a rank-revealing QR factorization of a matrix. We have developed a new algorithm based on Chandrasekaran and Ipsen's algorithm, with two advantages over it: First, our algorithm is much faster since we modify the main loop to accelerate its convergence and avoid the useless steps. Second, we apply a technique, suggested by Pan and Tang, that ensures termination, achieves the desired bounds, and fits into our theoretical studies.

Key words. orthogonal factorization, pivoting, rank-revealing QR factorization, numerical rank, serial algorithm, singular values.

AMS (MOS) subject classification. 15A03, 15A23, 15A42, 65F20, 65F25, 65F35.

1 Introduction

The rank-revealing QR factorization (RRQR factorization) is a valuable tool in numerical linear algebra because it detects the numerical rank of a matrix and because it provides the information needed about the rank and numerical nullspace to solve many rank-deficient linear least-squares problems. The RRQR factorization takes advantage of the efficiency and simplicity of the QR factorization, yet it produces information that is almost as reliable as that computed by means of the more expensive singular value decomposition.

We briefly summarize the properties of a rank-revealing QR factorization.

*All authors were partially supported by the European ESPRIT Project 9072-GEPPCOM.

[†]Departamento de Informática, Universidad Jaime I, Campus Penyeta Roja, 12071 Castellón, Spain, gquintan@inf.uji.es. This work was partially carried out while the author was a research fellow of the Department of Education and Science of the Valencian Government, and while visiting the Mathematics and Computer Science Division at Argonne National Laboratory.

[‡]Same address as first author; quintana@inf.uji.es.

Let A be an $m \times n$ matrix A (w.l.o.g. $m \geq n$) with singular values

$$\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_n \geq 0, \quad (1)$$

and define the numerical rank r of A with respect to the threshold τ as follows:

$$\sigma_r > \tau \geq \sigma_{r+1}.$$

Also, let A have a QR factorization of the form

$$AP = QR = Q \begin{pmatrix} R_{11} & R_{12} \\ 0 & R_{22} \end{pmatrix}, \quad (2)$$

where P is a permutation matrix, Q has orthonormal columns, R is upper triangular, and R_{11} is of order r . Further, let $\kappa(A)$ denote the two-norm condition number of a matrix A . We then say that (2) is an RRQR factorization of A if the following properties are satisfied:

$$\kappa(R_{11}) \approx \sigma_1/\sigma_r \quad \text{and} \quad \|R_{22}\|_2 = \sigma_{\max}(R_{22}) \approx \sigma_{r+1}. \quad (3)$$

Whenever a well-determined gap exists in the singular value spectrum between σ_r and σ_{r+1} , and hence the numerical rank r is well defined, the RRQR factorization (2) reveals the numerical rank of A by consisting of a well-conditioned leading submatrix R_{11} and a trailing submatrix R_{22} of small norm.

A list of problems where the RRQR factorization can be applied, as well as a complete overview of the classical algorithms for computing it, can be found in [8]. Here we will focus on the most recent studies.

In 1992 Hong and Pan proved that an optimal RRQR factorization is able to produce estimates of two consecutive singular values that are accurate up to a factor proportional to the matrix size [23]. This result is important because it shows that the RRQR may reliably reveal the numerical rank of any matrix. Unfortunately, Hong and Pan did not present a method for constructing such a factorization.

Shortly thereafter, two theoretic approximations were developed: the first by Chandrasekaran and Ipsen [14], the second one by Pan and Tang [24]. The main inconvenience of the former is its seriality, whereas the latter shows a higher degree of parallelism. On the other hand, Chandrasekaran and Ipsen's algorithm reaches tighter bounds than does the algorithm of Pan and Tang.

Specifically, Chandrasekaran and Ipsen proved that their algorithm halts and achieves the following bounds:

$$\begin{aligned} \sigma_{\min}(R_{11}) &\geq \frac{\sigma_k(A)}{\sqrt{k(n-k+1)}}, \\ \sigma_{\max}(R_{22}) &\leq \sigma_{k+1}(A) \sqrt{(k+1)(n-k)}, \end{aligned}$$

where $A \in \mathbb{R}^{m \times n}$ is the studied matrix and k is an integer value such that $1 \leq k < n$. In their experimental analysis, however, Chandrasekaran and Ipsen

found some termination problem caused by round-off arithmetic. That is, for some matrix types and sizes, the algorithm loops indefinitely and never halts. To solve this termination problem, Chandrasekaran and Ipsen proposed a solution based on a modification of the pivoting technique. The original pivoting technique assigns to each column a weight and then pivots the column with largest weight to the current position. The modified technique pivots the column with largest weight only if its weight is larger than the weight of the current column plus a constant. The proposed value for the constant is $n^2\epsilon$, where ϵ is the computer precision.

Pan and Tang used a different method to solve the termination problem. They proved that their most efficient algorithm, Algorithm 3 (reverse cyclic pivoting), achieves the following bounds:

$$\begin{aligned}\sigma_{\min}(R_{11}) &\geq \frac{f}{\sqrt{k(n-k+1)}}\sigma_k(A), \\ \sigma_{\max}(R_{22}) &\leq \frac{\sqrt{(k+1)(n-k)}}{f}\sigma_{k+1}(A),\end{aligned}$$

where f is a tolerance parameter such that $0 < f \leq 1/\sqrt{k+1}$. They also added a parameter M to take into account the errors introduced by the estimators. However, in practice since these are accurate, we have excluded that value.

We present here a variant of Chandrasekaran and Ipsen's algorithm that implements Pan and Tang's method for solving the termination problem. Our algorithm is faster than Chandrasekaran and Ipsen's [14], halts, achieves the desired bounds, and fits into our theoretical studies.

This paper is structured as follows. In Section 2, we review the original algorithm of Chandrasekaran and Ipsen and then present our new algorithm. In Section 3, we analyze the new algorithm theoretically. Finally, in Section 4, we summarize our results.

2 Algorithms with Guaranteed Termination

We start by briefly describing Chandrasekaran and Ipsen's algorithm, called Hybrid-III (see Figure 1). Figures 2 and 3 present two subalgorithms, Golub-I and Stewart-II, used by Hybrid-III. The former subalgorithm is based on the QR factorization with column pivoting [10]. The latter is based on Gragg and Stewart's algorithm [19]. Chandrasekaran and Ipsen proposed in their work the use of suffix -I for those algorithms that work on the lower right blocks of R , and suffix -II for those that work on the upper left block. The suffix -III is added to the algorithms that generate an RRQR and, therefore, work on the whole matrix. Both the Golub-I and Stewart-II subalgorithms carry out a unique permutation: the first one is restricted to block $R(k:n, k:n)$, the second to block $R(1:k, 1:k)$.

Algorithm Hybrid-III(k)
repeat
 repeat
 Golub-I(k)
 Stewart-II(k)
 until there are no more permutations.
 repeat
 Golub-I(k+1)
 Stewart-II(k+1)
 until there are no more permutations.
until there are no more permutations.
End Algorithm

Figure 1: Original Hybrid-III algorithm of Chandrasekaran Ipsen

Algorithm Golub-I(k)
1. Find the smallest index j , $k \leq j \leq n$, such that

$$\|R(k:j, j)\|_2 = \max_{k \leq i \leq n} \|R(k:i, i)\|_2$$
 .
2. **if** ($j > k$) **then**
3. Pivot column j to position k .
4. Retriangularize the matrix from the left with orthogonal transformations.
5. **end if**
End Algorithm

Figure 2: Golub-I subalgorithm

Algorithm Stewart-II(k)
1. Find the largest index j , $1 \leq j \leq k$, such that

$$\|e_j^T (R(1:k, 1:k))^{-1}\|_2 = \max_{1 \leq i \leq k} \|e_i^T (R(1:k, 1:k))^{-1}\|_2$$
 .
2. **if** ($j < k$) **then**
3. Pivot column j to position k .
4. Retriangularize the matrix from the left with orthogonal transformations.
5. **end if**
End Algorithm

Figure 3: Stewart-II subalgorithm

Algorithm Hybrid(GC)-III-sf(f, k)
repeat
 Golub-I-sf($f, k+1$)
 Golub-I-sf(f, k)
 Chan-II-sf(f, k)
 Chan-II-sf($f, k+1$)
until there are no more permutations.
End Algorithm

Figure 4: New algorithm Hybrid(GC)-III-sf

Since the technique used to solve the termination problem by Chandrasekaran and Ipsen was not included in their theoretical study, the effect on the reached bounds is unknown. Moreover, since it is based on a fixed value, independent from the matrix norm, it may be too small or large for some matrices.

We present in this paper a better solution to the termination problem, which is also included in our theoretical study. Specifically, we have solved the termination problem by means of a technique similar to that of Pan and Tang [24]. The solution is based on the use of a tolerance parameter f , $0 < f \leq 1$. The pivoting technique of the subalgorithms (Golub-I, etc.) assigns a weight to each column and then selects that column having the greatest weight as the pivot column. Our new pivoting technique consists of pivoting the heaviest column only if its weight scaled by the parameter f is larger than the weight of the current column.

Figures 4, 5, and 6 respectively present the new algorithm Hybrid-III(GC)-sf and its subalgorithms Golub-I-sf and Chan-II-sf. The new algorithm satisfies the following bounds (to be proved shortly):

$$\begin{aligned} \sigma_{\min}(R_{11}) &\geq \frac{f^2}{\sqrt{k(n-k+1)}} \sigma_k(A), \\ \sigma_{\max}(R_{22}) &\leq \frac{\sqrt{(k+1)(n-k)}}{f^2} \sigma_{k+1}(A), \end{aligned}$$

where f is a threshold interval such that $0 < f \leq 1$. This value is better than that of Pan and Tang ($0 < f \leq 1/\sqrt{k+1}$) because it allows tighter bounds.

Pan and Tang's algorithms as well as our new algorithm solve the termination problem. If the value of the parameter f is very small, the gap between the singular values is reduced, and the algorithm may fail to reveal the rank. On the other hand, a large value of f ($f = 1$ or very close to it) though adequate to reveal the rank, may produce termination problems. Pan and Tang propose

Algorithm Golub-I-sf(f,k)

1. Find the smallest index j , $k \leq j \leq n$, such that
$$\|R(k:j, j)\|_2 = \max_{k \leq i \leq n} \|R(k:i, i)\|_2 .$$
 2. **if** ($f \cdot \|R(k:j, j)\|_2 > |R(k, k)|$) **then**
 3. Move column j to position k .
 4. Retriangularize the matrix from the left with orthogonal transformations.
 5. **end if**
- End Algorithm**

Figure 5: “f-factor” variant of Golub-I subalgorithm

Algorithm Chan-II-sf(f,k)

1. Estimate the right singular vector associated with the smallest singular value of
$$R(1:k, 1:k).$$
 Store it in v .
 2. **if** ($f \cdot |v_j| > |v_k|$) **then**
 3. Pivot column j to position k .
 4. Retriangularize the matrix from the left with orthogonal transformations.
 5. **end if**
- End Algorithm**

Figure 6: “f-factor” variant of Chan-II subalgorithm

in their work the value $f = 0.65$, which seems to balance both situations.

Our new algorithm is faster than Hybrid-III because of the application of the following two techniques:

Use of Estimators: Stewart-II requires the computation of the inverse of an $r \times r$ upper triangular matrix, where r is the numerical rank. Since this process is very expensive (cubic), we use the subalgorithm Chan-II, which can be efficiently implemented if estimators are used. Specifically, we use the incremental condition estimator developed by Bischof [7], thus obtaining a quadratic cost, much smaller than the former. In theory, the substitution of Chan-II allows the possibility of a failure in the RRQR, that is, the rank will be revealed only if the estimates are accurate. Nevertheless, in practice, these estimators work correctly; indeed, for some of them, no matrix appears to make them fail.

Main-Loop Reorganization: The main loop of our new algorithm, clearly different from Chandrasekaran and Ipsen's algorithm, gives a faster convergence rate by avoiding the final steps that do not improve the column ordering.

Recently, Bischof and Quintana [8, 25] developed and implemented a new block algorithm for computing the RRQR factorization based on our method. Their experimental study, carried out on platforms including the IBM RS/6000-370, DEC AXP 6000-300, SGI R8000, HP 9000/715, and SUN hypersparc, has shown that the new algorithm based on our method is up to two or three times faster than the traditional QR factorization with column pivoting.

3 Theoretical Analysis of the New Algorithm

Next, we analyze the correctness of the new algorithm. First, we prove that it satisfies the desired bounds. Then, we prove that it halts.

Let k be an integer value $1 \leq k < n$, let f be a value $0 < f \leq 1$, and let $R \in \mathbb{R}^{n \times n}$ be a matrix such that $\sigma_k(R) > 0$ and $M\Pi = QR$ for some permutation matrix $\Pi \in \mathbb{R}^{n \times n}$ and $M \in \mathbb{R}^{m \times n}$. Then algorithm Hybrid(GC)-III-sf(f, k) achieves the following bounds:

$$\begin{aligned} \sigma_{\min}(R_{11}) &\geq \frac{f^2}{\sqrt{k(n-k+1)}} \sigma_k(M), \\ \sigma_{\max}(R_{22}) &\leq \frac{\sqrt{(k+1)(n-k)}}{f^2} \sigma_{k+1}(M). \end{aligned}$$

3.1 Background

Here we present some properties and a proposition required for the theoretical analysis. We present the proof only of the last proposition because the demonstration of the properties is widely-known.

Property 1 Let A be an $m \times n$ matrix and $\eta(A) = \max_{1 \leq j \leq n} \|Ae_j\|_2$. Then,

$$\eta(A) \leq \sigma_{\max}(A) \leq \sqrt{n}\eta(A).$$

Property 2 Consider a nonsingular $n \times n$ matrix A , and define $\tau(A) = 1/\max_{1 \leq i \leq n} \|e_i^T A^{-1}\|_2$. Then,

$$\frac{\tau(A)}{\sqrt{n}} \leq \sigma_{\min}(A) \leq \tau(A).$$

Property 3 Let $M\Pi = QR$ be the QR factorization of matrix M multiplied by some permutation matrix Π . Let R_{11} be the $k \times k$ upper left block of R , and let \bar{R}_{11} be the $(k+1) \times (k+1)$ upper left block of R . Then the following properties relate the singular values:

$$\begin{aligned} \sigma_k(M) &\geq \sigma_{\min}(R_{11}) \\ \sigma_{k+1}(M) &\geq \sigma_{\min}(\bar{R}_{11}). \end{aligned}$$

Property 4 Let $M\Pi = QR$ be the QR factorization of matrix M multiplied by a permutation matrix Π . Let R_{22} be the $(n-k) \times (n-k)$ lower right block of R , and let \bar{R}_{22} be the $(n-k+1) \times (n-k+1)$ lower right block of R . Then the following properties relate the singular values:

$$\begin{aligned} \sigma_{k+1}(M) &\leq \sigma_{\max}(R_{22}) \\ \sigma_k(M) &\leq \sigma_{\max}(\bar{R}_{22}). \end{aligned}$$

Proposition 5 Let A be an $m \times n$ matrix, and let $x \in \mathbf{R}^n$ satisfy $\|x\|_2 = 1$ and $\|Ax\|_2 = \epsilon$. Consider a permutation matrix Π such that $y = \Pi^T x$ satisfies $|y_n| \geq f \cdot \|y\|_\infty$, $0 < f \leq 1$. Then, if $A\Pi = QR$ is the QR factorization of matrix $A\Pi$,

$$|r_{nn}| \leq \frac{\sqrt{n}}{f} \epsilon.$$

Proof. Since $\Pi^T x = y$, $\|y\|_2 = \|x\|_2 = 1$. A well-known inequality between norms is $\|y\|_2 \leq \sqrt{n}\|y\|_\infty$.

From the hypothesis of the proposition, it is satisfied that $|y_n| \geq f/\sqrt{n}$.

On the other hand, it is possible to show that

$$Q^T Ax = Q^T A\Pi\Pi^T x = Ry = \begin{pmatrix} \vdots \\ r_{nn}y_n \end{pmatrix}.$$

Therefore,

$$\epsilon = \|Ax\|_2 = \|Q^T Ax\|_2 = \|Ry\|_2 \geq |r_{nn}y_n| \geq |r_{nn}| \frac{f}{\sqrt{n}},$$

and the desired result is obtained

$$|r_{nn}| \leq \frac{\sqrt{n}}{f} \epsilon.$$

QED.

3.2 Desired Bounds

Let k be an integer value $1 \leq k < n$, let f be a value $0 < f \leq 1$, and let $R \in \mathbb{R}^{n \times n}$ be a matrix such that $\sigma_k(R) > 0$ and $M\Pi = QR$ for some permutation matrix $\Pi \in \mathbb{R}^{n \times n}$ and $M \in \mathbb{R}^{m \times n}$. Then algorithm Hybrid(GC)-III-sf(f,k) achieves the following bounds:

$$\begin{aligned} \sigma_{\min}(R_{11}) &\geq \frac{f^2}{\sqrt{k(n-k+1)}} \sigma_k(M) \\ \sigma_{\max}(R_{22}) &\leq \frac{\sqrt{(k+1)(n-k)}}{f^2} \sigma_{k+1}(M). \end{aligned}$$

Proof. First we briefly describe the notation that will be used in the demonstration. The upper triangular matrix R is partitioned as follows:

$$R = \begin{pmatrix} \bar{A} & \bar{B} \\ & \bar{C} \end{pmatrix} = \begin{pmatrix} A & B \\ & C \end{pmatrix} = \begin{pmatrix} \hat{A} & \hat{B} \\ & \hat{C} \end{pmatrix},$$

where \bar{A} is $(k-1) \times (k-1)$, \bar{C} is $(n-k+1) \times (n-k+1)$, $A = R_{11}$ is $k \times k$, $C = R_{22}$ is $(n-k) \times (n-k)$, \hat{A} is $(k+1) \times (k+1)$, and \hat{C} is $(n-k-1) \times (n-k-1)$.

Next we prove that the algorithm reaches the preceding bounds whenever it halts. In such cases, it is obvious that there are no permutations in any of the four subalgorithms. Let's check what bounds are achieved if there are no permutations.

First, Golub-I-sf(f,k+1) is analyzed. If it does not permute any column, then $\|Ce_1\|_2 \geq f\|Ce_j\|_2$, $1 \leq j \leq n-k$.

Therefore,

$$|r_{k+1,k+1}| = \|Ce_1\|_2 \geq f \max_{1 \leq i \leq n-k} \|Ce_i\|_2.$$

Applying property 1 to matrix C , we obtain

$$\sigma_{\max}(C) \leq \sqrt{n-k} \max_{1 \leq i \leq n-k} \|Ce_i\|_2,$$

and combining both expressions we have

$$\sigma_{\max}(C) \leq \frac{\sqrt{n-k}}{f} |r_{k+1,k+1}|. \quad (4)$$

Second, Golub-I-sf(f,k) is analyzed. If it does not permute any column, then $\|\bar{C}e_1\|_2 \geq f\|\bar{C}e_j\|_2, 1 \leq j \leq n - k + 1$.

Therefore,

$$|r_{kk}| = \|\bar{C}e_1\|_2 \geq f \max_{1 \leq i \leq n-k+1} \|\bar{C}e_i\|_2.$$

Applying property 1 to matrix \bar{C} , we obtain

$$\sigma_{\max}(\bar{C}) \leq \sqrt{n-k+1} \max_{1 \leq i \leq n-k+1} \|\bar{C}e_i\|_2,$$

and combining both expressions we have

$$\sigma_{\max}(\bar{C}) \leq \frac{\sqrt{n-k+1}}{f} |r_{kk}|. \quad (5)$$

Third, algorithm Chan-II-sf(f,k) is analyzed. If it does not permute any column, from proposition 5 the following condition is satisfied

$$|r_{kk}| \leq \sigma_{\min}(A) \frac{\sqrt{k}}{f}.$$

If this expression is combined with the inequality (5), the following is obtained:

$$\frac{\sqrt{k}}{f} \sigma_{\min}(A) \geq |r_{kk}| \geq \frac{f}{\sqrt{n-k+1}} \sigma_{\max}(\bar{C}),$$

and from it

$$\sigma_{\min}(R_{11}) \geq \frac{f^2}{\sqrt{k(n-k+1)}} \sigma_{\max}(\bar{C}).$$

Applying property 4, we obtain the first required bound:

$$\sigma_{\min}(R_{11}) \geq \frac{f^2}{\sqrt{k(n-k+1)}} \sigma_k(M).$$

Finally, Chan-II-sf(f,k+1) is analyzed. If no permutations occur, from proposition 5 the following condition is satisfied

$$|r_{k+1,k+1}| \leq \sigma_{\min}(\hat{A}) \frac{\sqrt{k+1}}{f}.$$

If this expression is combined with the inequality (4), the following is obtained:

$$\frac{\sqrt{k+1}}{f} \sigma_{\min}(\hat{A}) \geq |r_{k+1,k+1}| \geq \frac{f}{\sqrt{n-k}} \sigma_{\max}(C),$$

and from it

$$\frac{\sqrt{(n-k)(k+1)}}{f^2} \sigma_{\min}(\hat{A}) \geq \sigma_{\max}(R_{22}).$$

Applying property 3, we obtain the second required bound:

$$\frac{\sqrt{(n-k)(k+1)}}{f^2} \sigma_{k+1}(M) \geq \sigma_{\max}(R_{22}).$$

QED.

3.3 Termination

The second step of our analysis is to prove that the algorithm halts. If no permutations occur, the algorithm terminates with a unique iteration of the main loop.

If there are permutations, however, the termination of the algorithm must be demonstrated. The basic goal is to prove that $|\det(R_{11})|$ is a monotonically increasing function during the execution of the algorithm. Since the value $|\det(R_{11})|$ is unique for each different ordering of the columns of R , no two orderings of the columns will be equal if $|\det(R_{11})|$ increases continuously. Since a finite number of different orderings exist for the columns of R , the algorithm must eventually terminate.

In each iteration of the algorithm, four calls are performed, say, Golub-I-sf(f, k), Golub-I-sf($f, k+1$), Chan-II-sf(f, k), and Chan-II-sf($f, k+1$). Subalgorithms Golub-I-sf($f, k+1$) and Chan-II-sf(f, k) do not modify $|\det(R_{11})|$. However, Golub-I-sf(f, k) and Chan-II-sf($f, k+1$) do modify it. It must be proved that after the execution of Golub-I-sf(f, k) and Chan-II-sf($f, k+1$), the value of $|\det(R_{11})|$ is larger.

First, the consequences of Golub-I-sf(f, k) in $|\det(R_{11})|$ are analyzed. Assume that a permutation occurred. Then column j , $k \leq j \leq n$, must be moved to position k . The permutation can be performed in two steps:

1. Swap $k+1$ and j .
2. Swap k and $k+1$.

The first step does not modify $|\det(R_{11})|$, whereas the second step does. Consider that the first step has been carried out, and focus on the second step.

Let \bar{R} be the triangular matrix before the permutation, and let \bar{R}_{11} be the $k \times k$ upper left block of \bar{R} . Let R be the triangular matrix after the permutation has been applied, and let R_{11} be the $k \times k$ upper left block of R . Since we must demonstrate that the determinant is always increasing, it is sufficient to prove that $|\det(R_{11})| > |\det(\bar{R}_{11})|$.

Consider that the matrix before the permutation has the following structure:

$$\bar{R} = \begin{pmatrix} x & x & \dots & x & x & \dots & x \\ & x & \dots & x & x & \dots & x \\ & & & \vdots & \vdots & & \vdots \\ & & & \alpha & \beta & \dots & x \\ & & & & \gamma & \dots & x \\ & & & & & & \vdots \\ & & & & & & x \end{pmatrix} \begin{matrix} k \\ k+1 \end{matrix} .$$

Then, $|\det(\bar{R}_{11})| = |p\alpha|$, where p is the product of the first $k-1$ diagonal elements. Now columns k and $k+1$ are permuted and the matrix is retriangularized. Since there has been a permutation, $|\alpha| < f\sqrt{\beta^2 + \gamma^2}$.

After the permutation and retriangularization, the matrix has the following structure:

$$R = \begin{pmatrix} x & x & \dots & x & x & \dots & x \\ & x & \dots & x & x & \dots & x \\ & & & \vdots & \vdots & & \vdots \\ & & & \sqrt{\beta^2 + \gamma^2} & x_1 & \dots & x \\ & & & & x_2 & \dots & x \\ & & & & & & \vdots \\ & & & & & & x \end{pmatrix} \begin{matrix} k \\ k+1 \end{matrix} ,$$

where $\sqrt{x_1^2 + x_2^2} = |\alpha|$. The value of $|\det(R_{11})|$ is

$$|\det(R_{11})| = |p\sqrt{\beta^2 + \gamma^2}| > |p| \frac{|\alpha|}{f} = \frac{1}{f} |p\alpha| \geq |\det(\bar{R}_{11})| ,$$

and therefore the desired conclusion is reached:

$$|\det(R_{11})| > |\det(\bar{R}_{11})| .$$

Next, the consequences of Chan-II-sf(f,k+1) in $|\det(R_{11})|$ are analyzed. Assume that a permutation occurred. Then column j must be moved to position $k+1$. The permutation can be performed in two steps:

1. Swap j and k .
2. Swap k and $k+1$.

The first step does not change $|\det(R_{11})|$, whereas the second step may modify $|\det(R_{11})|$. Consider that the first step has been carried out and focus on the second step.

Let \bar{R} be the triangular matrix before the permutation, and let \bar{R}_{11} be the $k \times k$ upper left block of \bar{R} . Let R be the triangular matrix after the permutation, and let R_{11} be the $k \times k$ upper left block of R . Since we must prove that the value of the determinant always increases, it is sufficient to prove that $|\det(R_{11})| > |\det(\bar{R}_{11})|$.

Consider that the matrix before the permutation has the following structure:

$$\bar{R} = \begin{pmatrix} x & x & \dots & x & x & \dots & x \\ & x & \dots & x & x & \dots & x \\ & & & \vdots & \vdots & & \vdots \\ & & & \alpha & \beta & \dots & x \\ & & & & \gamma & \dots & x \\ & & & & & & \vdots \\ & & & & & & x \end{pmatrix} \begin{matrix} k \\ k+1 \end{matrix} .$$

Then, $|\det(\bar{R}_{11})| = |q\alpha|$, where q is the product of the first $k-1$ diagonal elements.

Now columns k and $k+1$ of the matrix must be permuted and the matrix retriangularized. Since a permutation has occurred, then

$$|\gamma| > \sigma_{\min}(\bar{R}(1:k+1, 1:k+1)) \frac{\sqrt{k+1}}{f}. \quad (6)$$

After the permutation and the retriangularization, the matrix has the following structure:

$$R = \begin{pmatrix} x & x & \dots & x & x & \dots & x \\ & x & \dots & x & x & \dots & x \\ & & & \vdots & \vdots & & \vdots \\ & & & \sqrt{\beta^2 + \gamma^2} & x_1 & \dots & x \\ & & & & x_2 & \dots & x \\ & & & & & & \vdots \\ & & & & & & x \end{pmatrix} \begin{matrix} k \\ k+1 \end{matrix} ,$$

where $\sqrt{x_1^2 + x_2^2} = |\alpha|$. The new value of $|\det(R_{11})|$ is

$$|\det(R_{11})| = |q\sqrt{\beta^2 + \gamma^2}| .$$

After the application of Chan-II-sf(f,k+1), the following condition is satisfied:

$$|x_2| \leq \sigma_{\min}(R(1:k+1, 1:k+1)) \frac{\sqrt{k+1}}{f}. \quad (7)$$

Combining the inequalities (6) and (7) and since $\sigma_{\min}(\bar{R}(1:k+1, 1:k+1)) = \sigma_{\min}(R(1:k+1, 1:k+1))$, it is obtained that

$$|\gamma| > |x_2|.$$

Assume that the retriangularization is carried out by means of a Givens rotation. Then

$$\begin{pmatrix} c & s \\ -s & c \end{pmatrix} \begin{pmatrix} \beta \\ \gamma \end{pmatrix} = \begin{pmatrix} \sqrt{\beta^2 + \gamma^2} \\ 0 \end{pmatrix},$$

and

$$\begin{pmatrix} c & s \\ -s & c \end{pmatrix} \begin{pmatrix} \alpha \\ 0 \end{pmatrix} = \begin{pmatrix} x_1 \\ x_2 \end{pmatrix}.$$

If c and s are isolated from the first two equations and their expressions are substituted in the fourth equation, then

$$\frac{x_2}{\alpha}\beta + \frac{x_1}{\alpha}\gamma = 0,$$

which is equivalent to

$$x_2\beta + x_1\gamma = 0.$$

Applying some of the properties of the modulus $|\cdot|$ to the preceding expression, we obtain

$$\begin{aligned} |x_2\beta| &= |x_1\gamma| \\ |x_2||\beta| &= |x_1||\gamma|. \end{aligned}$$

Since $|\gamma| > |x_2|$, then $|x_2||\beta| > |x_1||x_2|$, and we obtain from this expression that

$$|\beta| > |x_1|.$$

Thus,

$$\begin{aligned} |\det(R_{11})| &= |q\sqrt{\beta^2 + \gamma^2}| \\ &> |q\sqrt{x_1^2 + x_2^2}| \\ &\geq |\alpha|. \end{aligned}$$

Therefore, the required result is obtained:

$$|\det(R_{11})| > |\det(\bar{R}_{11})|.$$

QED.

4 Concluding Remarks

In this paper we presented a new algorithm for computing the rank-revealing QR factorization (RRQR) of triangular matrices. Chandrasekaran and Ipsen's algorithm has a termination problem with some matrix types and sizes. To solve this problem, these authors proposed a method that was not included in their theoretical studies. Hence, the effect of this change on the algorithm bounds is unknown. In contrast, we solve the termination problem with a method first suggested by Pan and Tang that has been fully included in our new theoretical study. In addition, the new algorithm is faster because

- Subalgorithm Chan-II, with a quadratic cost on the rank, is used instead of Stewart-II, with a cubic cost.
- The main structure of the algorithm has been modified in order to allow a faster convergence rate.

We also theoretically analyze the new algorithm to prove that it halts and achieves the desired bounds.

Experimental study has shown that the new implementation runs between 2 times (for low-rank matrices) and 40 times (for high-rank matrices) faster than Chandrasekaran and Ipsen's algorithm on 200×200 matrices. We expect that this factor will grow higher for larger matrices.

Because of the advantages of the new algorithm, it has been used by Bischof and Quintana [8] to develop a new block algorithm for computing the RRQR factorization. This block algorithm has shown a much higher performance than the traditional QR factorization with column pivoting. In an experimental study carried out on several different platforms (including an IBM RS/6000-370, DEC AXP 6000-300, HP 9000/715, SUN hypersparc, and SGI R8000), the new block algorithm based on our method has been about two to three times faster than QR factorization with column pivoting.

Acknowledgments

We express our gratitude to Dr. C. H. Bischof for his interesting suggestions, ideas, and discussions. We also thank Gail Pieper for her effort to improve the original manuscript.

References

- [1] E. Anderson, Z. Bai, C. Bischof, J. Demmel, J. Dongarra, J. DuCroz, A. Greenbaum, S. Hammarling, A. McKenney, S. Ostrouchov, and D. Sorensen. *LAPACK User's Guide*. SIAM, Philadelphia, 1992.
- [2] Christian H. Bischof. A block QR factorization algorithm using restricted pivoting. In *Proceedings SUPERCOMPUTING '89*, pages 248–256, Baltimore, Md., 1989. ACM Press.

- [3] Christian H. Bischof. Incremental condition estimation. *SIAM Journal on Matrix Analysis and Applications*, 11(2):312–322, 1990.
- [4] Christian H. Bischof. A parallel QR factorization algorithm with controlled local pivoting. *SIAM Journal on Scientific and Statistical Computing*, 12(1):36–57, 1991.
- [5] Christian H. Bischof and Per Christian Hansen. Structure-preserving and rank-revealing QR factorizations. *SIAM Journal on Scientific and Statistical Computing*, 12(6):1332–1350, November 1991.
- [6] Christian H. Bischof and Per Christian Hansen. A block algorithm for computing rank-revealing QR factorizations. *Numerical Algorithms*, 2(3-4):371–392, 1992.
- [7] Christian H. Bischof and Ping Tak Peter Tang. Robust incremental condition estimation. Preprint MCS-P225-0391, Mathematics and Computer Science Division, Argonne National Laboratory, 1991.
- [8] Christian H. Bischof and Gregorio Quintana-Ortí. Computing Rank-Revealing QR factorizations of Dense Matrices. Preprint MCS-P559-0196, Mathematics and Computer Science Division, Argonne National Laboratory, 1996.
- [9] Christian H. Bischof and Gregorio Quintana-Ortí. Codes for Rank-Revealing QR factorizations of Dense Matrices. Preprint MCS-P560-0196, Mathematics and Computer Science Division, Argonne National Laboratory, 1996.
- [10] P. A. Businger and G. H. Golub. Linear least squares solution by Householder transformation. *Numerische Mathematik*, 7:269–276, 1965.
- [11] Tony F. Chan. Rank revealing QR factorizations. *Linear Algebra and Its Applications*, 88/89:67–82, 1987.
- [12] Tony F. Chan and Per Christian Hansen. Some applications of the rank revealing QR factorization. *SIAM Journal on Scientific and Statistical Computing*, 13(3):727–741, 1992.
- [13] Tony F. Chan and Per Christian Hansen. Low-rank revealing QR factorizations. *Numerical Linear Algebra and Applications*, 1(1):33–44, 1994.
- [14] Shivkumar Chandrasekaran and Ilse Ipsen. On rank-revealing QR factorizations. *SIAM Journal on Matrix Analysis and Applications*, 15(2), 1994.
- [15] J. J. Dongarra, J. R. Bunch, C. B. Moler, and G. W. Stewart. *LINPACK Users' Guide*. SIAM Press, Philadelphia, 1979.

- [16] L. V. Foster. Rank and null space calculations using matrix decomposition without column interchanges. *Linear Algebra and Its Applications*, 74:47–71, 1986.
- [17] G. H. Golub, V. Klema, and G. W. Stewart. Rank degeneracy and least squares problems. Technical Report TR-456, University of Maryland, Dept. of Computer Science, 1976.
- [18] Gene H. Golub and Charles F. Van Loan. *Matrix Computations*. The Johns Hopkins University Press, 1983.
- [19] W. B. Gragg and G. W. Stewart. A stable variant of the secant method for solving nonlinear equations. *SIAM Journal on Numerical Analysis*, 13(6):889–903, 1976.
- [20] Ming Gu and Stanley Eisenstat. A stable and efficient algorithm for the rank-one modification of the symmetric eigenproblem. Technical Report YALEU/DCS/RR-916, Yale University, Department of Computer Science, 1992.
- [21] Per Christian Hansen, Takishi Sekii, and Hiromoto Shibahashi. The modified truncated SVD-method for regularization in general form. *SIAM Journal on Scientific Computing*, 13:1142–1150, 1992.
- [22] Nicholas J. Higham. Efficient algorithms for computing the condition number of a tridiagonal matrix. *SIAM Journal on Scientific and Statistical Computing*, 7:150–165, 1986.
- [23] Yoo Pyo Hong and Chiang-Tsuan Pan. The rank revealing QR decomposition and SVD. *Mathematics of Computation*, 58:213–232, 1992.
- [24] Chiang-Tsuan Pan and Ping Tak Peter Tang. Bounds on singular values revealed by QR factorization. Technical Report MCS-P332-1092, Mathematics and Computer Science Division, Argonne National Laboratory, 1992.
- [25] Gregorio Quintana-Ortí. *Algoritmos secuenciales, orientados por bloques y paralelos para el cálculo del rango numérico matricial*. Ph.D. thesis, Universidad Politécnica de Valencia, May 1995.
- [26] Gregorio Quintana-Ortí, Xiaobai Sun, and Christian H. Bischof. A BLAS-3 version of the QR factorization with column pivoting. Preprint MCS-P551-1295, Mathematics and Computer Science Division, Argonne National Laboratory, 1995.
- [27] Lothar Reichel and William Gragg. Fortran subroutines for updating the QR factorization. *ACM Transactions on Mathematical Software*, 16:369–377, 1990.

- [28] Robert Schreiber and Charles Van Loan. A storage efficient WY representation for products of Householder transformations. *SIAM Journal on Scientific and Statistical Computing*, 10(1):53–57, 1989.
- [29] G. W. Stewart. Rank degeneracy. *SIAM Journal on Scientific and Statistical Computing*, 5:403–413, 1984.
- [30] G. W. Stewart. An updating algorithm for subspace tracking. Technical Report CS-TR-2494, University of Maryland, Department of Computer Science, 1990.
- [31] G. W. Stewart. Determining rank in the presence of error. Technical Report CS-TR-2972, Dept. of Computer Science, University of Maryland, 1993.